MNRAS **000**, 000–000 (0000) Preprint 11 October 2016 Compiled using MNRAS LATEX style file v3.0

Stellar Velocity Dispersion in Mergers: The Effects of Dust and Star Formation

Nathaniel R. Stickley* and Gabriela Canalizo†
Department of Physics and Astronomy, University of California, 900 University Ave, Riverside, CA 92521 USA

11 October 2016

arXiv:1610.02601v1 [astro-ph.GA] 9 Oct 2016

ABSTRACT
We investigate the effects of stellar evolution and dust on measurements of stellar velocity dispersion in mergers of disk galaxies. N-body simulations and radiative transfer analysis software are used to obtain mass-weighted and flux-weighted measurements of stellar velocity dispersion. We find that the distribution of dust with respect to the distribution of young stars in such systems is more important than the total degree of attenuation. The presence of dust typically causes flux-weighted measurements of stellar velocity dispersion to be elevated with respect to mass-weighted measurements because dust preferentially obscures young stars, which tend to be dynamically cooler than older stellar populations in such systems. In exceptional situations, in which young stars are not preferentially obscured by dust, flux-weighted velocity dispersion measurements tend to be negatively offset with respect to mass-weighted measurements because the dynamically cool young stellar populations are more luminous, per unit mass, than older stellar populations. Our findings provide a context for comparing observationally-obtained measurements of velocity dispersion with measurements of velocity dispersion obtained from galaxy merger simulations.

Key words: methods:numerical – galaxies:kinematics and dynamics – ISM:dust – ISM:extinction

1 INTRODUCTION

Central stellar velocity dispersion (σ_*) is an important observable quantity in galactic astronomy because it correlates with many other galactic properties, such as depth of the galaxy's potential well, the mass of its central supermassive black hole (Ferrarese & Merritt 2000; Gebhardt et al. 2000; Tremaine et al. 2002), its luminosity, and its surface brightness distribution (Davies et al. 1987; Dressler et al. 1987; Bender et al. 1992).

Although σ_* has been measured from observations and in numerical simulations for decades, little work has been done to compare the fundamentally different methods that are used to measure σ_* in real galaxies versus simulated galaxies. Thus, comparing values of σ_* in the galaxy simulation literature with observationally-obtained measurements of σ_* has been problematic.

Observationally, σ_* can only be measured by analysing the light emitted by a galaxy. This is an inherently flux-weighted measurement method. In contrast, the natural way of measuring σ_* in a numerical simulation involves using the velocity and mass data in the simulation to compute a mass-weighted velocity dispersion. This mass-weighted measurement technique has been used for as long as σ_* has been measured in simulations(e.g., Villumsen 1982; Cox et al. 2006; Robertson et al. 2006; Johansson et al. 2009). If real galaxies consisted of stars with a uniform mass-to-light ratio (Υ) and contained no gas nor dust, observations of σ_* should, in principle, agree perfectly with the quantity computed from simulation data. However, reality is not so simple; galaxies typically contain a broad variety of stars and at least a small amount of dust.

In our earlier work (Stickley & Canalizo 2012, hereafter denoted SC1), we found that including a toy model for dust attenuation—a uniformly dense slab of attenuating material—could potentially cause a decrease in the flux-weighted velocity dispersion ($f\sigma_*$), relative to the mass-weighted velocity dispersion ($m\sigma_*$) in a simulated galaxy. Similar results were obtained analytically by Baes & Dejonghe (2000), who studied the effect of diffusely distributed dust on measurements of stellar kinematics in elliptical galaxies. Unlike our toy model and the smooth distributions of dust assumed by Baes & Dejonghe (2000), real interstellar dust is not smoothly distributed; it occurs in clumps, sheets, and filaments as well as in a more diffuse state with varying density. Thus, based on previous work, it is not clear how a realistic dust distribution would affect the value of $f\sigma_*$.

* E-mail: nstic001@ucr.edu
† E-mail: gabriela.canalizo@ucr.edu

© 0000 The Authors

Even in the absence of dust, comparisons of $f\sigma_*$ and $m\sigma_*$ are not straightforward because stars in galaxies are not uniformly luminous. In a previous paper (Stickley & Canalizo 2014, hereafter denoted SC2), we found that measurements of $m\sigma_*$ that were based upon newly-formed stars in simulated galaxies were significantly lower than measurements of $m\sigma_*$ that included stars of all ages. This finding was consistent with the observations of the so-called "σ_* discrepancy" (Rothberg & Fischer 2010; Rothberg et al. 2013), in which different stellar populations in a single galaxy yield discrepant values if σ_*. Since young stellar populations are more luminous per unit mass (i.e., have smaller Υ), the presence of young stars may weight the value of $f\sigma_*$ downward relative to $m\sigma_*$.

In this paper, we examine the effects of dust and stellar evolution on measurements of $f\sigma_*$ in order to gain insights into the differences between the value of σ_* measured observationally and the value of σ_* reported in the galaxy simulation literature. We describe our methods in Section 2 and present our findings in Section 3. General conclusions and limitations to the work are discussed in Section 4.

2 METHODOLOGY

We began by selecting a small set of snapshots from a binary galaxy merger simulation that was previously analysed in SC2. Using these snapshots, we generated synthetic Doppler-broadened galaxy spectra. We analysed these synthetic spectra to obtain flux-weighted velocity dispersion measurements. Then, the particle mass and velocity information in the snapshots was analysed to obtain mass-weighted velocity dispersion measurements. A method of consistently comparing the mass-based and flux-based measurement methods was developed so that intrinsic, systematic effects could be separated from the effects of dust and star formation.

2.1 Numerical Simulations

We have used simulation snapshots from merger S1, described in SC2. Since the details of this simulation have already been discuss in detail, we will only summarize the key features here; refer to SC2 for the full details of the simulation.

The simulation was a binary, 1:1 mass ratio, prograde-prograde merger of disk galaxies. It was performed using the N-body, SPH code, GADGET-3 (Springel 2005). Each progenitor consisted of 1.6×10^6 particles. Eighty percent of the disk mass of each progenitor galaxy initially consisted of stars with the other 20% in the form of gas and dust. The gravitational softening length of stars was 25 pc, meaning that physical processes on scales smaller than 25 pc could not be resolved. Each stellar particle in the simulation represented a population of stars, rather than individual stars. The interstellar medium was modelled as a multi-phase gas with hot and cool phases. The hot gas was able to cool radiatively, while cool gas was able become heated by stellar feedback and AGN feedback. A sub-resolution approximation was used to include stellar formation in the simulation; new stellar particles were spawned in cool, dense regions of gas at a rate designed to match observational evidence. As stars evolved and added metals to the ISM, the metallicity of the gas increased. Each stellar particle that formed during the simulation carried a variable specifying its creation time and metallicity. The latter was set equal to the metallicity of the gas from which the particle was spawned.

2.2 Mass-Weighted Velocity Dispersion

The mass-weighted velocity dispersion ($m\sigma_*$) was computed with GSnap[1] (N.R. Stickley, in preparation), using the same technique described in SC1 and SC2. In summary, a virtual rectangular slit of width $w = 2$ kpc and length $\ell = 20$ kpc was placed on the galaxy of interest. A viewing direction, (θ, ϕ) and slit position angle α, were then specified. The masses and velocities of all stars appearing in the slit were used to compute $m\sigma_*$, according to

$$m\sigma_* = \sqrt{v_i^2 m_i / M - (v_i m_i / M)^2} \tag{1}$$

with

$$M = \sum_i m_i$$

where the standard summation convention has been utilized; repeated indices imply a sum over that index. Using this technique allowed us to directly compare the present work with the results presented in SC2. The primary source of uncertainty in this measurement was particle noise, which never exceeded 0.9% of the measured value. Note that this measurement technique can confuse velocity dispersion with rotation.

2.3 Flux-Weighted Velocity Dispersion

In this section, we describe the processes that were used to obtain and analyse the synthetic, Doppler-broadened, galaxy spectra.

2.3.1 Simulated Spectra

Synthetic, Doppler-broadened galaxy spectra were generated using the polychromatic, Monte Carlo, radiative transfer code, Sunrise (Jonsson 2006; Jonsson et al. 2010; Jonsson & Primack 2010). Sunrise was capable of creating realistic images of galaxies from arbitrary viewing directions. More importantly, the code could compute a high resolution spectrum for each pixel of each image that it generated.

Given a GADGET-3 snapshot file, Sunrise began by discretising the spatial domain of the simulation using an adaptive mesh. A Monte Carlo radiative transfer algorithm was then performed on the discretised volume. The algorithm assumed that the dust content of the ISM was proportional to the metallicity of the SPH particles in the simulation snapshot file. In this work, we assumed that 40% of the metals in the ISM occurred in the form of dust. Because the stellar particles in our galaxy simulation represented entire stellar populations, rather than individual stars, each stellar particle in Sunrise emitted a spectrum corresponding to a population of stars. These spectra were pre-computed using

[1] http://www.gsnap.org

STARBURST99 (Leitherer et al. 1999). Since our GADGET-3 simulations only tracked the ages and metallicities of stellar particles that formed during the simulation, we manually assigned ages and metallicities to the pre-existing stellar particles; all particles were assumed to form instantaneously, 6 Gyr before the beginning of the simulation with a metallicity of $Z = 0.025$, which is approximately 1.25 $Z_{\odot}$. SUNRISE treated stellar particles younger than 10 Myr as active star-forming regions, containing enhanced dust concentrations and a photodissociation region. These star-forming regions are modelled using the code, MAPPINGSIII (Dopita et al. 2005; Groves et al. 2008). AGN emission was not included in the radiative transfer computation.

Using the stellar spectra and the spatial distribution of stars and dust, SUNRISE computed the resulting spectra by tracing 10^7 "photon bundles" through the simulation domain. Each photon bundle carried spectral data on a random walk through the simulated galaxy and accounted for Doppler shifts due to the motion of the individual stellar particles. The bundles were eventually collected by a virtual integral field spectrograph. As the photons were collected, the Doppler shifts of the individual bundles combined to form Doppler-broadened spectra. In addition to computing the dust-attenuated spectral information for each pixel of the generated image, SUNRISE computed the unattenuated spectra. This allowed us to clearly identify the effect of attenuation in each virtual observation, by comparing attenuated and unattenuated fluxes.

Besides the galaxy simulation snapshot file and the input spectra, the other main user input to the code was the choice of dust model (i.e., the dust grain size distribution and albedo). For this work, we used the Milky Way dust model from Weingartner & Draine (2001). It should be noted that the dust model used by SUNRISE significantly affects the final spectrum of the simulated galaxy (Jonsson et al. 2010).

The radiative transfer computation required a large amount of memory, partially owing to the fact each photon bundle carried a detailed spectrum. In order to reduce the memory requirement and accelerate the computation, we limited our spectral coverage to the Mg Ib region, from 5040 Å to 5430 Å. Our input stellar spectra contained 1170 wavelength bins in this region, logarithmically spaced. Once SUNRISE had finished generating synthetic spectra for the individual pixels of each image, we placed a rectangular slit, measuring 2×20 kpc, on the image. The spectra of all pixels appearing within the slit were then combined to form a single spectrum. This combined spectrum was analysed in order to determine the flux-weighted velocity dispersion ($f\sigma_*$).

2.3.2 *Measuring Velocity Dispersion from Spectra*

Stellar velocity information is encoded in all galaxy spectra, since the light emitted by a galaxy consists of the sum of the Doppler shifted spectra of its constituent stars. Similarly, each of our synthetic spectra consisted of a sum of Doppler-shifted particle spectra. In order to decode the spectra and recover the velocity information, we used the penalized pixel-fitting code, pPXF, developed by Cappellari (Cappellari & Emsellem 2004; Cappellari 2012). In general, the pPXF algorithm worked by fitting a parametrized model spectrum, $G_{\rm mod}(x)$, to an observed galaxy spectrum, $G(x)$. One of the parameters of $G_{\rm mod}(x)$ was $f\sigma_*$. Thus, $f\sigma_*$ was ultimately determined by finding the model spectrum that best fit each galaxy spectrum. In our analysis, we made no attempt to correct for the effect of galaxy rotation.

More specifically, the model spectrum was a linear combination of template spectra convolved with a parametrized line of sight velocity distribution (LOSVD):

$$G_{\rm mod}(x) = \sum_{k=1}^{K} w_k [B * T_k](x) + \sum_{l=0}^{L} b_l \mathcal{P}_l(x) \qquad (w_k \geq 0), \quad (2)$$

where w_k are weights, $B(x) = \mathcal{L}(cx)$ is a broadening function, with $\mathcal{L}(v)$ the LOSVD, c is the speed of light, T_k is a library of template spectra, and $*$ denotes convolution. A linear combination of Legendre polynomials, $\mathcal{P}_l(x)$ (with weights b_l) was used to account for low-frequency differences between the shape of the templates and the shape of the galaxy spectrum.

The LOSVD was expanded as a Gauss-Hermite series,

$$\mathcal{L}(v) = \frac{\exp(-y^2/2)}{\sigma\sqrt{2\pi}} \left[1 + \sum_{m=3}^{M} h_m H_m(y)\right], \quad (3)$$

where H_m are Hermite polynomials, $y \equiv (v - V)/\sigma$, and $(V, \sigma, h_3, h_4, \ldots, h_M)$ are free parameters related to the moments of the velocity distribution. For example, V is the mean line-of-sight velocity, σ corresponds to the standard deviation (i.e., $f\sigma_*$), h_3 is related to the skewness, and h_4 is related to the kurtosis. Optimal values for these parameters, as well as the weights, w_k and b_l, were found using a non-linear least-squares minimization algorithm.

The quantity minimized by the least-squares optimization routine was the objective function:

$$\chi^2_{\rm p} = \chi^2(1 + \lambda^2 \mathcal{D}^2), \quad (4)$$

where χ^2 is given by

$$\chi^2 = \sum_{n=1}^{N} \left[\frac{G_{\rm mod}(x_n) - G(x_n)}{\Delta G(x_n)}\right]^2, \quad (5)$$

with $\Delta G(x_n)$ the measurement error on $G(x_n)$. The $\mathcal{D}^2$ in Eq. (4) is a penalty term, given by the integrated square deviation of $\mathcal{L}(v)$ from its best-fitting Gaussian, $\mathcal{G}(v)$,

$$\mathcal{D}^2 = \frac{\int_{-\infty}^{\infty} [\mathcal{L}(v) - \mathcal{G}(v)]^2 {\rm d}v}{\int_{-\infty}^{\infty} \mathcal{G}^2(v) {\rm d}v} \quad (6)$$

The penalty term clearly increases as $\mathcal{L}(v)$ deviates from a pure Gaussian. Thus, it has the effect of forcing the fitting routine to favour LOSVDs that are more nearly Gaussian. The parameter, λ, was a user-specified quantity that allowed us to adjust the importance of the penalty term. Setting $\lambda = 0$ caused the best-fitting LOSVD to be a general Gauss-Hermite series. Increasing the parameter λ resulted in LOSVDs that were more nearly Gaussian. When using pPXF with noisy spectra, obtained observationally, non-zero values of λ are often used in order to force the fitting routine to favour Gaussian LOSVDs. This helps pPXF to partially ignore the effect of noise in the spectra in cases for which the true LOSVD is Gaussian (Cappellari & Emsellem 2004).

Our spectra did not include the sort of noise that is present in real (i.e., observed) spectra, but we adjusted λ nonetheless in order to examine the full range of velocity dispersions that could be obtained using our spectra. Specifically, we set $\lambda = 0$ to determine the value of $f\sigma_*$ that would be measured by a researcher who prefers to approximate their LOSVDs using a Gauss-Hermite series. We also set $\lambda = 10^6$ in order to determine the value of $f\sigma_*$ that would be reported by a researcher who prefers to use simple Gaussian LOSVDs.

We constructed a template library using a subset of the raw spectra emitted by the stellar particles in the SUNRISE computation. The template library was gradually expanded to include more spectra until the effect of adding additional spectra no longer significantly effected the χ^2 of the resulting fits. In total, 67 template spectra were included in the library. For the measurement error, $\Delta G(x_n)$, we assumed a uniform value for each wavelength bin in the spectrum. Specifically, we set each entry equal to 1% of the mean signal strength (i.e., $\Delta G(x_n) = 0.01\langle G\rangle$, for all n). This had no effect on the resulting values $f\sigma_*$; it merely determined the magnitude of the formal measurement uncertainty.

2.4 Consistent Comparisons

In order to identify the effects of interstellar dust and star formation on measurements of $f\sigma_*$, we first needed a general method for consistently comparing measurements of $m\sigma_*$ with measurements of $f\sigma_*$. For instance, both measurement techniques needed to include exactly the same region of space, measurement uncertainties needed to be quantified, and intrinsic differences between the two measurement methods needed to be identified.

2.4.1 Slit Calibration and Antialiasing

A virtual observatory was positioned 5 Mpc from the centre of the simulated system. Each observation was performed using a virtual integral field spectrograph with a resolution of 400×400 pixels. The field of view of each array was 200 kpc, thus each pixel represented a 0.5×0.5 kpc region of the simulation. In order to compare the mass-weighted and flux-weighted velocity dispersions, we first verified that we were measuring the same region of the simulation snapshot using our two methods. The verification process involved comparing detailed images, generated by SUNRISE, with images generated using GSNAP's interactive particle visualization feature to confirm that the same particles were included in both types of measurements.

Once the location and orientation of the slits had been calibrated, there remained a mismatch between the sizes of the flux slit and the mass slit, due to the pixelated nature of the flux data. Pixels falling on the border of the slit typically extended outside of the slit. Thus, flux from a large region outside of the slit contributed to the measurement of $f\sigma_*$. To reduce this aliasing effect, we assigned weights to the pixels, depending on the degree of overlap with the slit. The result can be seen in Figure 1, where white corresponds to a weight of zero and black corresponds to a weight of unity. Edge pixels were weighted intermediately, so they appear as shades of grey. These weights were applied to the pixel spectra when the combined slit spectrum was computed. Furthermore, we varied the size of the mass slit in order to determine the maximum and minimum possible values of $m\sigma_*$ in the region that was measured by the flux-weighting method. The blue rectangles in Figure 1 indicate the smallest and largest slit sizes that were used during this process. In Section 3, we report the value of $m\sigma_*$ measured using the fiducial 2×20 kpc slit. We use the maximum and minimum values of $m\sigma_*$ over the full range of slit sizes as upper and lower bounds on the uncertainty whenever $m\sigma_*$ is compared with $f\sigma_*$.

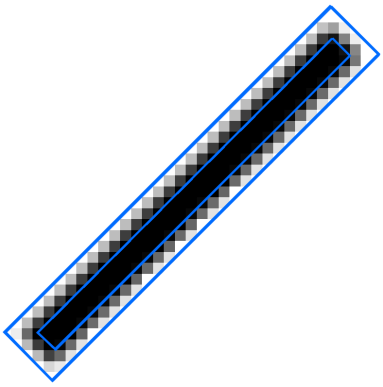

Figure 1. An illustration of the pixel-weighting scheme. White pixels indicate weights of zero while black pixels indicate weights equal to unity. Gray pixels on the edge of the slit are weighted intermediately. The blue rectangles indicate the smallest and largest slit sizes that were used when measuring $m\sigma_*$. The fiducial slit (not shown) lies mid-way between the blue rectangles.

2.4.2 An Intrinsic Measurement Discrepancy

In order to characterize the intrinsic differences between the two measurement methods, we compared measurements of $m\sigma_*$ and $f\sigma_*$ in snapshots consisting of identical stars (i.e., stars of a specific mass, age, metallicity, and Υ) and containing no dust nor gas. The positions and velocities of the stars in these snapshots were identical to the snapshots that were later analysed. In such systems, each star particle contributed equally to the measurement of $m\sigma_*$ because the stellar masses were equal. All stars emitted identical spectra and had identical luminosities. Thus, the inverse square variation of flux with distance and the random noise introduced by the Monte Carlo radiative transfer scheme used by SUNRISE were the only effects that prevented each star from contributing equally to the $f\sigma_*$ measurement (in the absence of random noise, the stars nearest to the camera would have contributed approximately 4% more flux than the stars farthest from the camera).

Upon comparing the techniques, we found a systematic discrepancy. The flux-weighted measurement technique yielded values that were elevated with respect to the fiducial mass-weighted value by an average of 2.7% when a Gauss-Hermite LOSVD was assumed and 3.3% when a Gaussian LOSVD was assumed. After accounting for measurement uncertainty due to slit mismatch, these offsets fell to 1.4% and 2.0%, respectively. The maximum offset from the fiducial mass-weighted value was 5.8% (Gauss-Hermite) and 6.2% (Gaussian). No negative offsets were observed; *flux-based*

measurements always agreed with or slightly exceeded the mass-based measurements.

We do not believe that this discrepancy is a fundamental, real effect. It may have been caused by an aspect of the spectral template library creation process. It also may have been due to a peculiarity of the pPXF fitting code or the SUNRISE radiative transfer code. Note that, while it was not *caused* by the pixel-weighting scheme, a better pixel-weighting scheme may have been able to reduce the discrepancy. Regardless of the exact cause, we were able to take the effect into account when performing our later analysis because the effect was small and bounded in our set of observations. Any measurement of $f\sigma_*$ that exceeded the corresponding $m\sigma_*$ measurement by more than 6.2% was attributable to the effects of dust attenuation or non-uniform Υ. Any negative offsets larger than the slit mismatch uncertainty could also be attributed to these effects.

2.5 Sample Selection

In general, our observations were chosen to include a mixture of extreme and ordinary situations. More precisely, snapshots that were known to exhibit extreme values of $m\sigma_*$ were included along with snapshots that were known to be in a dynamically relaxed state. Lines of sight known to have enhanced dust extinction (for instance, parallel to a disk of gas) were included along with lines of sight with far less dust extinction (e.g., nearly perpendicular to disk structures). Intermediate cases were included as well.

A total of seven snapshots from the simulation were examined using our mass-weighted and flux-weighted velocity dispersion measurement methods. The first six of these were the snapshots labelled *c*–*h* in SC2. These correspond to interesting points in the dynamical evolution of the merger. The seventh snapshot was the final snapshot of the simulation. SUNRISE was used to perform virtual observations on each snapshot along three viewing directions. For each resulting image, we placed four slits—centred on the same pixel, but rotated uniformly about that pixel's centre. Thus, we performed 21 virtual observations and used 84 individual slits. The basic features of the snapshots are summarized in Figure 2. Detailed descriptions and renderings of each of the 21 observations are presented in the Appendix. To gain a better understanding of the detailed dynamics of the merger as a whole, refer to SC2.

3 RESULTS

3.1 Overview

The lower panel of Figure 2 summarizes all of the velocity dispersion measurements performed on snapshots 1–7. Measurements of $m\sigma_*$ are shown as black boxes. Measurements of $f\sigma_*$ that assumed a pure Gaussian LOSVD, are plotted as blue circles. Measurements of $f\sigma_*$ that used the more general Gauss-Hermite series to model the LOSVD are plotted as red diamonds. Filled circles and diamonds indicate that the flux-weighted measurement was based on a spectrum that included dust attenuation. Open circles and diamonds indicate that dust was ignored when computing the spectrum. Error bars are only included in the plot when the measurement uncertainty is larger than the plotted symbol (this plotting convention is used throughout this paper). The horizontal axis indicates the simulation time at which the snapshot was saved. Within each snapshot bin, there are three groups of four measurements, corresponding to the three viewing directions and four slit orientations per viewing direction. Proceeding from left to right, the first four measurements were made along the direction of Camera 1, the second four represent Camera 2 measurements, and the final four represent Camera 3 measurements (refer to the Appendix for more information about the camera positions). For context, the top panel of Figure 2 shows the evolution of the mass-weighted σ_* as a function of simulation time. The central panel indicates the separation distance between the supermassive black holes at the centres of the galaxies, which is a proxy for the nuclear separation distance.

In general, all measurement methods indicated that the velocity dispersion of the dynamically excited system in Snapshot 4 was elevated with respect to the values in snapshots 3, 5, 6, and 7. It was also clear that the systems in snapshots 5–7 were very similar to one another, which was expected, since the merger remnant was passively evolving when the final three snapshots were obtained.

Several trends are apparent in the data presented in Figure 2. For instance, most flux-weighted measurements that assumed a Gaussian LOSVD fell above the corresponding measurement that used a Gauss-Hermite series to model the LOSVD. It was also clear that, in most cases, the $f\sigma_*$ measurements that included dust attenuation fell above $m\sigma_*$, while the non-attenuated $f\sigma_*$ typically fell below $m\sigma_*$. These trends are examined in the following sections.

3.2 Gaussian versus Gauss-Hermite Fitting Methods

Since researchers use different methods to extract the value of $f\sigma_*$ from real spectral data, we have examined the differences between the two most common flux-based methods. As mentioned previously, these two methods differ due to the assumed functional form of the LOSVD. On one extreme, a Gaussian LOSVD is assumed, on the other, a Gauss-Hermite series is used to model the LOSVD. We will use the symbol, $\hat{f}\sigma_*$ to denote a measurement made using the Gaussian method and $\tilde{f}\sigma_*$ to denote a measurement based on the Gauss-Hermite method.

In Figure 3, we plot the fractional offset between $\hat{f}\sigma_*$ and $\tilde{f}\sigma_*$. When dust attenuation was not included in the spectra, there was no measurable, systematic offset between the two methods. However, when dust attenuation was included, $\hat{f}\sigma_*$ exceeded $\tilde{f}\sigma_*$ by 6.7%, on average. This offset was not seen in some of the measurements made on Snapshot 1 and Snapshot 3. Recall that these snapshots were recorded during the first pass and between the first and second passes of the merger, respectively.

3.3 Mass-based versus Flux-based Measurements

In Figures 4 and 5, we compare measurements of $m\sigma_*$ with measurements of $f\sigma_*$ in the set of observations, described in Section 2.5. For the measurements that included dust, the mean fractional offsets between flux-weighted and mass-

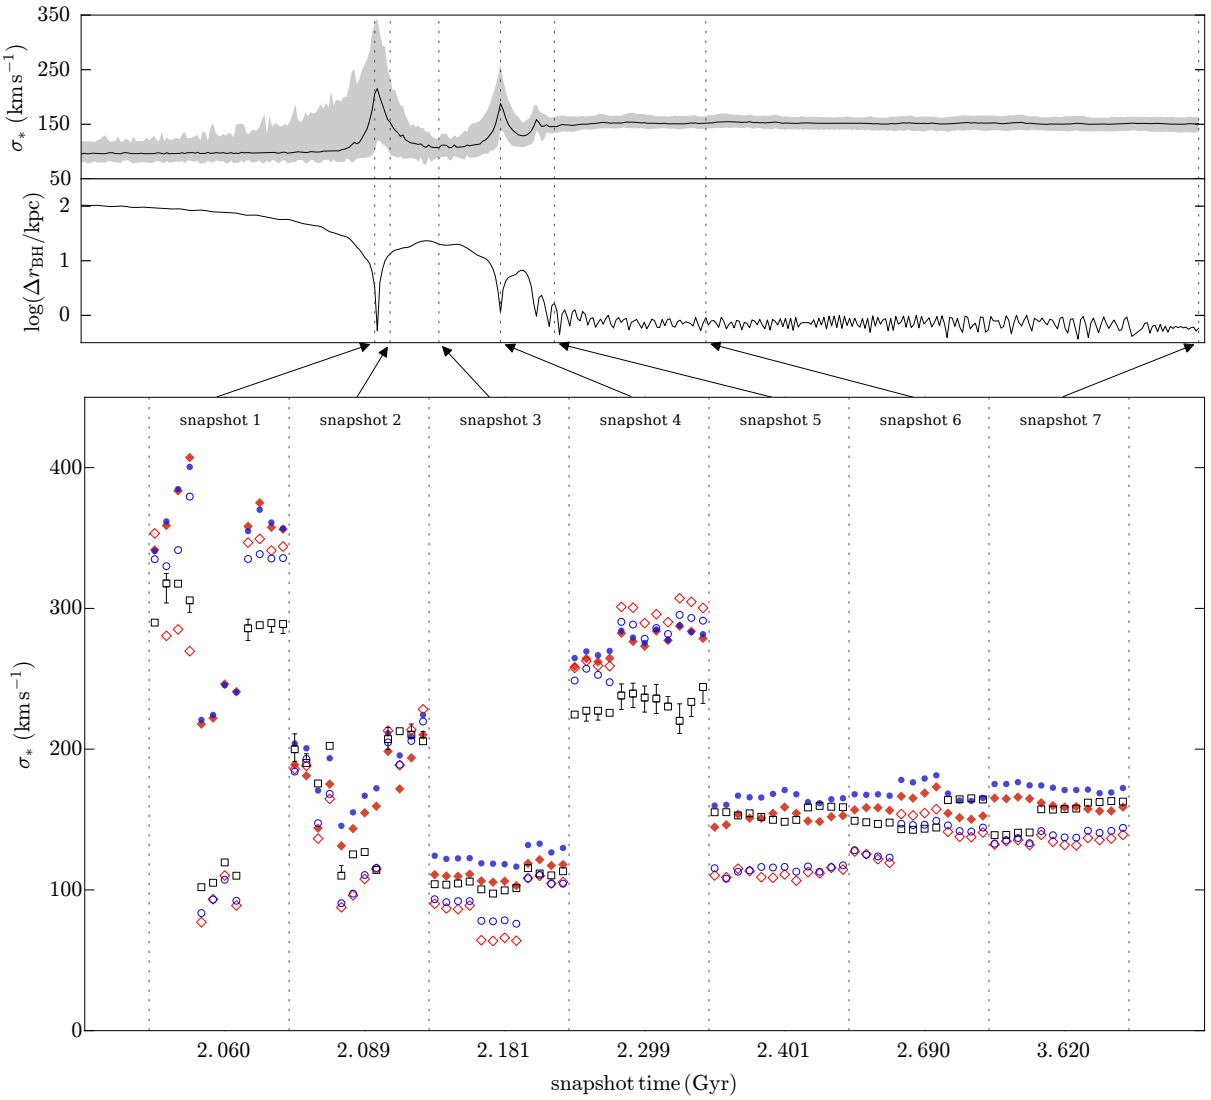

Figure 2. Upper panel: the mass-weighted σ_* as a function of simulation time. The dark line indicates the mean value of σ_*, measured along 1000 random lines of sight. The upper and lower bounds of the shaded region indicate the maximum and minimum values of σ_* along the set of 1000 directions. Middle panel: the separation distance between the supermassive black holes in the nuclei of the progenitor galaxies. Lower panel: An overview of all velocity dispersion measurements performed on snapshots 1–7. Refer to Section 3.1 for the meanings of the symbols used in the plot.

weighted velocity dispersions, $(f\sigma_* - m\sigma_*)/m\sigma_*$, were respectively 0.19 and 0.13 for $\hat{f}\sigma_*$ and $\tilde{f}\sigma_*$. The offsets were largest along the direction of Camera 2 in Snapshot 1. Interestingly, this particular camera direction was chosen, in part, because of the apparently small amount of intervening dust along its line of sight. The large discrepancy between $m\sigma_*$ and $f\sigma_*$ was likely related to the dynamically extreme nature of this snapshot. For measurements that neglected dust attenuation, the mean fractional offsets were respectively -0.06 and -0.07 for $\hat{f}\sigma_*$ and $\tilde{f}\sigma_*$.

The mean fractional offsets presented above were computed using a very broad variety of snapshots. If we limit the sample to dynamically passive (i.e., non-merging) systems by removing snapshots 1–4 from the analysis, the fractional offsets of $\hat{f}\sigma_*$ become 0.11 and -0.15 for dusty and dustless systems, respectively. The corresponding fractional offsets of $\tilde{f}\sigma_*$ become 0.05 and -0.16.

Many of the individual offsets exceeded the threshold set by the intrinsic measurement discrepancy. The mean offsets typically exceeded the threshold as well. Two general trends were observed:

- Dust attenuation often caused $f\sigma_*$ to be elevated with respect to $m\sigma_*$.
- When dust attenuation was neglected, $f\sigma_*$ tended to be smaller than $m\sigma_*$.

Note that many more snapshots from a larger variety of simulations would be needed in order to compute robust values for these offsets. Also note that these offsets were observed in a set of galaxies with ongoing star formation. Consequently, the trends in the offset only strictly apply to systems with ongoing or recent star formation. The trends are discussed further in the following sections.

The reader may have noticed that the positive offset

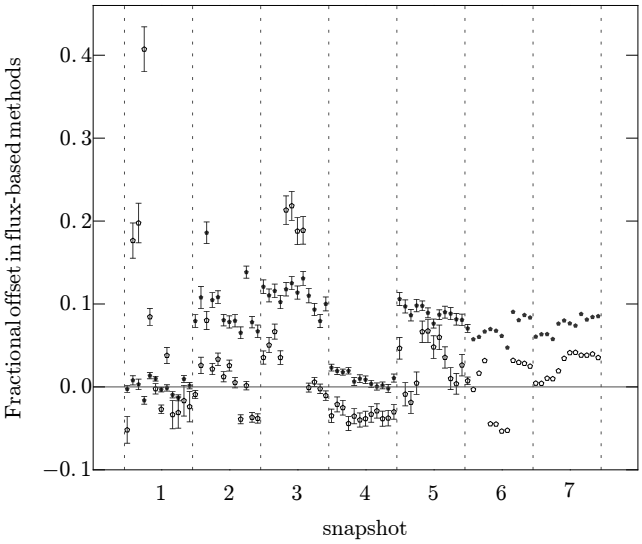

Figure 3. The quantity $(\hat{f}\sigma_* - \tilde{f}\sigma_*)/\tilde{f}\sigma_*$ for each measurement slit. This is the fractional offset between flux-weighted velocity dispersions measured under the assumption of a Gaussian LOSVD ($\hat{f}\sigma_*$) relative to those assuming a Gauss-Hermite LOSVD ($\tilde{f}\sigma_*$). When dust attenuation was included in the synthetic spectra (solid symbols), the Gaussian model consistently yielded larger velocity dispersions in all snapshots except for Snapshot 1 and Snapshot 3. When dust attenuation was not included (open symbols), there was no clear trend in the offset.

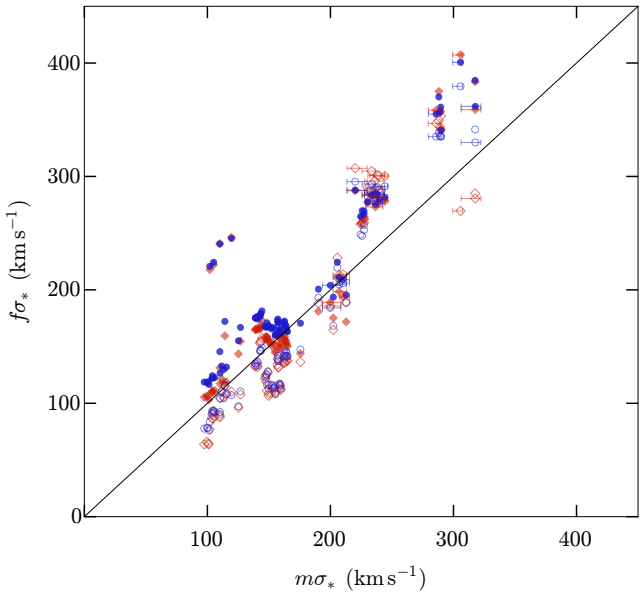

Figure 4. Flux-weighted versus mass-weighted velocity dispersion. The flux-weighted quantity tends to be positively offset with respect to the mass-weighted quantity when dust is included (solid circles and diamonds). The positive offset is accentuated when the LOSVD is assumed to be Gaussian (blue circles). On the other hand, when dust attenuation is not taken into account (open diamonds and circles), the flux-weighted quantity is negatively offset with respect to the mass-weighted quantity.

due to the presence of dust was smaller, on average, than the negative offset present when dust was ignored. This occurred even though the flux-weighted measurement technique had an intrinsic positive offset. We do not consider this to be a robustly-determined trend because the magnitude of the offset due to dust could easily depend upon the amount of gas and dust present in the simulation snapshot. The magnitude may also depend upon the type of dust present in the system or upon the resolution of the underlying simulation. Furthermore, the relative fraction of young stars likely also affects the magnitude of the offset. Using our limited sample, we can only determine very general trends, such as the *direction* of the offsets, with confidence.

3.4 Dust Attenuation

The analysis in Section 3.3 revealed that the presence of dust generally increased $f\sigma_*$, relative to $m\sigma_*$. From the results of SC2, we also know that younger stars in merger simulations exhibited lower $m\sigma_*$ than older stars, *on average*. Based on this evidence, it seems likely that dust *preferentially* obscured the dynamically cooler young stars. This effectively removed the dynamically cooler stars from the measurement of $f\sigma_*$ and thus increased the measured value. This conclusion is consistent with the well-known fact that stars tend to be born in dusty environments and is further supported by the snapshot renderings in the Appendix, which show that the dust and young stars tended to be concentrated in the same regions of the system.

In principle, dust can also increase $f\sigma_*$ measurements by scattering the light of high velocity stars into the line of sight of the camera, as described by Baes & Dejonghe (2002). This effect is likely weak in our observations because all of our measurement slits are relatively small and centred upon galactic nuclei, while the effect of scattering is most important in the outer regions of galaxies.

There were a few exceptions to the general trend, discussed above; some measurements of $f\sigma_*$ that included dust attenuation were negatively offset with respect to $m\sigma_*$. These exceptions likely occurred in cases for which the attenuation due to dust was weak or the dust did not attenuate the light of young stars preferentially. The attenuation of young stars may have been weak for three reasons:

(i) The total amount of dust in the system was small.
(ii) The dust and young stars were distributed in a flattened sheet (or disk) and the observation was made along a line of sight nearly perpendicular to the sheet.
(iii) The dust became decoupled from the young stars, due to a collision, heating, or winds.

In the first case, the total degree of dust attenuation would be low. In the second and third cases, the total degree of attenuation could be significant, but the light of young stars would not be *preferentially* attenuated.

Since we obtained dust-attenuated and dust-free images for each camera direction, we were able to determine the level of attenuation in each slit. We defined attenuation as,

$$\mathcal{A} = \log\left(\frac{F_{\rm tot}}{F_{\rm dust}}\right), \qquad (7)$$

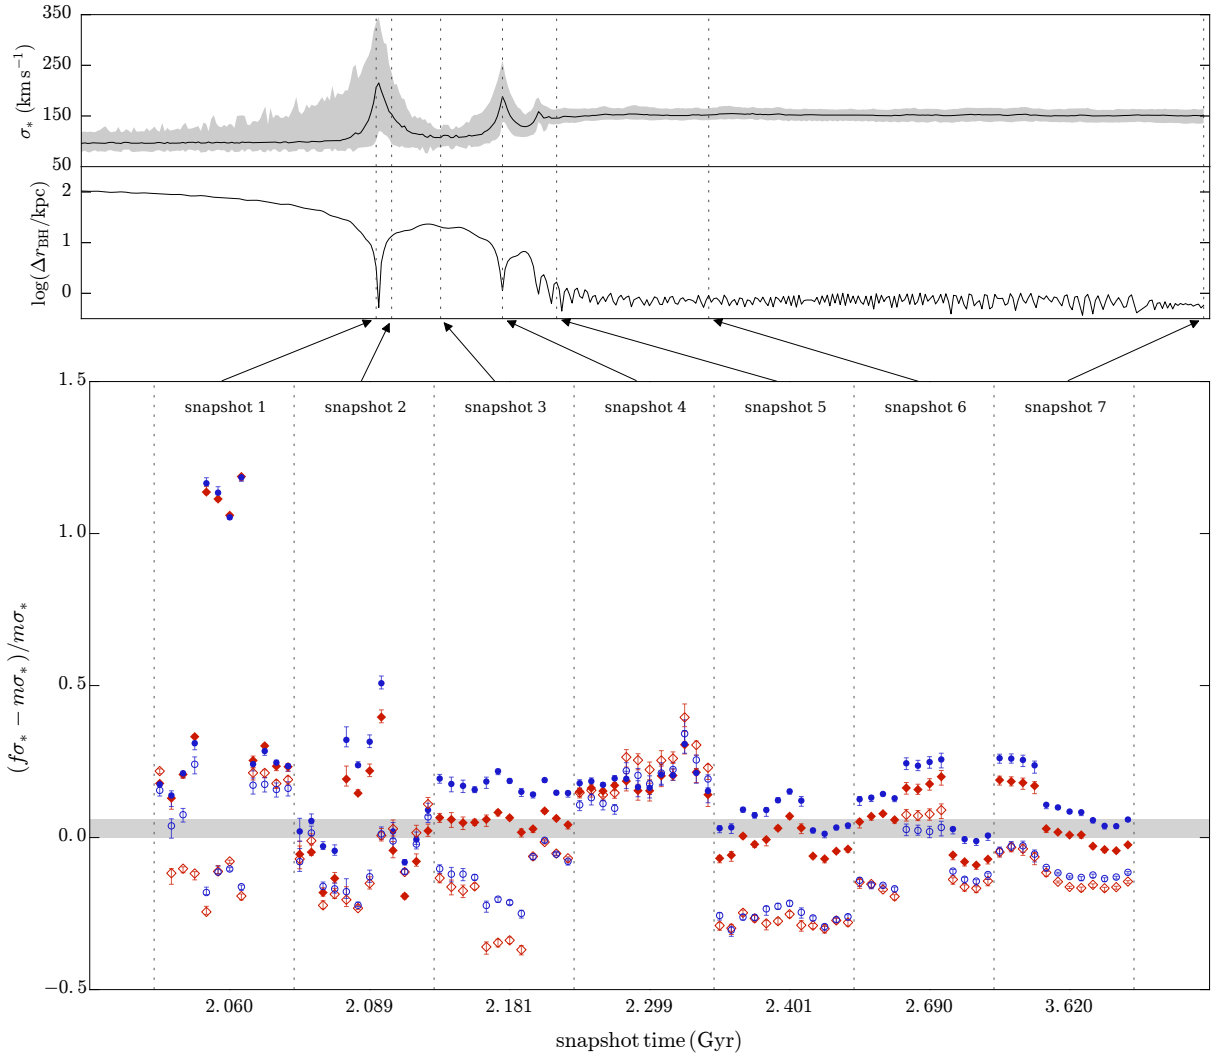

Figure 5. Upper panels: Refer to the caption of Figure 2. Lower panel: The fractional offset between flux-weighted and mass-weighted velocity dispersions. The shaded region indicates the range of offset values that may have been caused by the intrinsic discrepancy between the flux-weighted and mass-weighted measurement techniques. Only offsets falling outside of this region were considered significant. When dust was included, a significant positive offset was observed in a majority of the 84 measurement slits (63 $\hat{f}\sigma_*$ measurements and 43 $\tilde{f}\sigma_*$ measurements). In the absence of dust, the offset was negative in the majority of measurements (57 $\hat{f}\sigma_*$ measurements and 58 $\tilde{f}\sigma_*$ measurements).

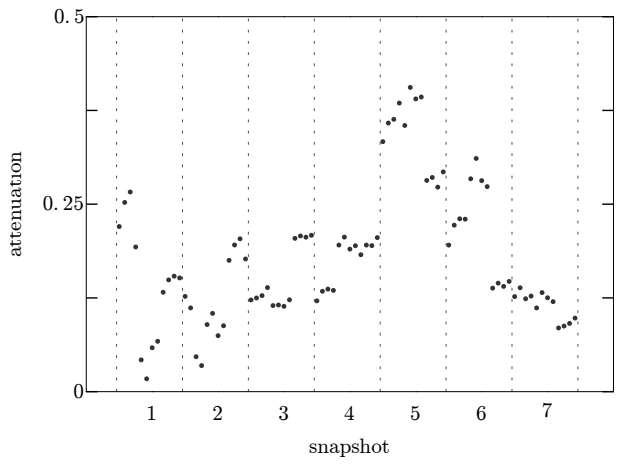

Figure 6. The attenuation, $\mathcal{A}$, for each measurement slit.

where $F_{\rm tot}$ is the total flux passing through the measurement slit when dust was ignored and $F_{\rm dust}$ is the flux received when dust was included in the SUNRISE analysis. The value of $\mathcal{A}$ for all measurement slits is presented in Figure 6. By comparing this figure with Figure 5, one can easily see that the total attenuation was not unusually low in situations for which the dust-attenuated measurement of $f\sigma_*$ was negatively offset, relative to $m\sigma_*$. Furthermore, when the fractional offset, $(f\sigma_* - m\sigma_*)/m\sigma_*$, is plotted as a function of attenuation, as in Figure 7, no clear relationship is evident; total attenuation did not correlate with the fractional offset of $f\sigma_*$ from $m\sigma_*$. We conclude that the distribution of dust with respect to the distribution of young stars was more important than the total amount of dust present in the simulations. In other words, the effect of dust on measurements of $f\sigma_*$ can not be determined by the amount of reddening present.

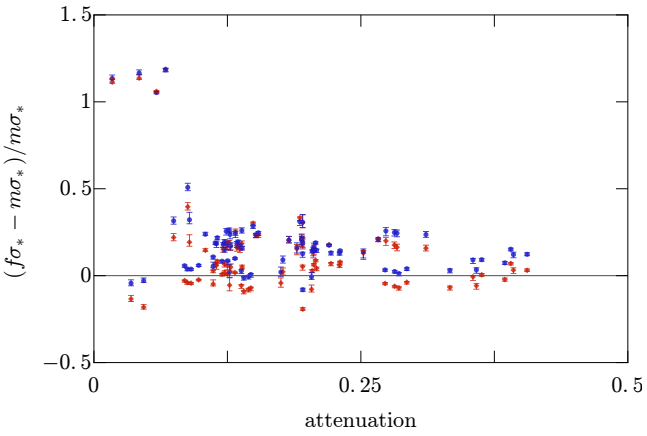

Figure 7. The fractional offset of $f\sigma_*$ with respect to $m\sigma_*$ versus attenuation, $\mathcal{A}$. It is clear that higher attenuation does not correlate with a larger offset.

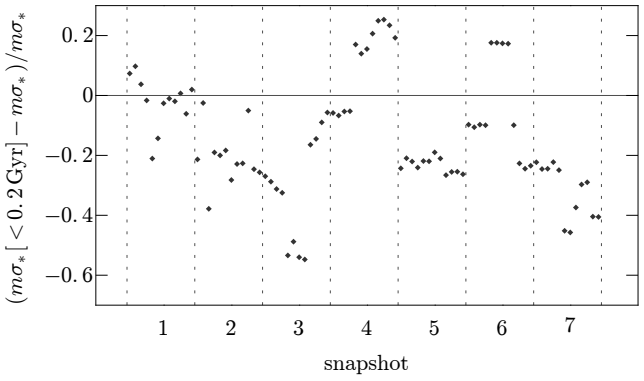

Figure 8. The fractional offset of $m\sigma_*$ for stars younger than 200 Myr, relative to the $m\sigma_*$ of *all* stars. The offset tends to be negative, indicating that the velocity dispersion of younger stars tends to be lower than that of older stars.

3.5 The $m\sigma_*$ of Young Stars

In this section, we investigate the exceptions to the second trend observed in Section 3.3. Specifically, these were cases in which the dust-free $f\sigma_*$ was positively offset with respect to $m\sigma_*$. In principle, the lower Υ (i.e., higher luminosity per unit mass) of the young stellar populations should cause any flux-weighted quantity to shift toward the characteristic value exhibited by the young populations (when dust attenuation is neglected). It is then reasonable to assume that the offsets observed in the dust-free $f\sigma_*$ measurements were due to offsets present in the population of young stars. We measured the mass-weighted velocity dispersion using only the relatively young stellar populations in order to determine whether the observed offsets were due to the dynamics of the young stars.

In Figure 8, we show the fractional offset of $m\sigma_*$ for stars less than 200 Myr old, relative to the global population. A positive offset was measured in 17 of the 84 individual measurements. Most of the positively-offset dust-free $f\sigma_*$ measurements in Figure 5 can be easily explained using the data presented in Figure 8; the measurements of $f\sigma_*$ were positively offset because the σ_* of young stars was elevated. Computing $m\sigma_*$ in an appropriate stellar age bin (i.e., something other than 0–200 Myr) would likely explain the remaining positive offsets.

While young stars were dynamically cooler than the global population, *on average*, individual measurements of $m\sigma_*$ in the young population were able to exceed the global value. This occurred in dynamically peculiar situations (e.g., in snapshots 1 and 4, which were galactic collisions) and also when young stars were in a rotating disk, viewed edge-on (as in the case of Camera 2 of snapshot 6).

4 DISCUSSION

By analysing snapshots from a simulation of merging disk galaxies, we were able to compare mass-weighted velocity dispersion measurements ($m\sigma_*$) with flux-weighted velocity dispersion measurements ($f\sigma_*$). The measurements of $f\sigma_*$ were performed by generating synthetic spectra with the radiative transfer code, SUNRISE. The spectra were then analysed using the pPXF code of Cappellari (2012)—a code which is commonly used to measure velocity dispersions from real (i.e., observationally-obtained) spectra. Two spectra were obtained for each observation slit—one that included the effect of dust attenuation and one that ignored the presence of dust. All of our simulated spectra included effects of stellar evolution. Our primary findings were:

(i) Dust preferentially obscured the light of young stars because young stars were often found in dusty environments. Therefore, dust partially removed the dynamically cool young stars from the measurement of $f\sigma_*$, while stars of all ages contributed equally to the $m\sigma_*$ measurement. This caused measurements of $f\sigma_*$ to be elevated with respect to $m\sigma_*$ in most cases.

(ii) When dust was ignored, measurements of $f\sigma_*$ tended to fall below their $m\sigma_*$ counterparts. This was due to the higher luminosity-to-mass ratio (i.e., smaller Υ) of the young stellar populations, which weighted $f\sigma_*$ toward the velocity dispersion of the dynamically cooler young stars.

In exceptional cases, the dust-free measurements of $f\sigma_*$ exceeded $m\sigma_*$ or the dust-attenuated measurements of $f\sigma_*$ fell below $m\sigma_*$. The former type of exception occurred when the dynamics of the young stars were peculiar or when the system was observed along a fortuitous line of sight that caused the velocity dispersion of the young population to appear elevated. The latter type of exception occurred when the dust did not obscure the young stellar population more significantly than it obscured the older population. We intentionally chose snapshots and viewing directions that we knew would increase the likelihood of finding such exceptional cases.

We also found that the total degree of attenuation due to dust was not a good predictor of the offset between $f\sigma_*$ and $m\sigma_*$. Observations with significant dust attenuation often exhibited a smaller offset due to dust than observations that suffered less attenuation. In other words, the distribution of the dust was more important than the total amount of dust present. Results of previous research indicated that diffusely-distributed dust can have the opposite effect; it can lower the value of $f\sigma_*$ relative to $m\sigma_*$. In SC1, we found that

the presence of a large attenuating slab resulted in measurements of $f\sigma_*$ that were lower than $m\sigma_*$. This happened because the stars of highest dispersion (in the centre of the galaxy) were more strongly attenuated. Baes & Dejonghe (2000) found similar results when analytically studying diffusely distributed dust in elliptical galaxies. Thus, dust may have opposing effects on measurements of $f\sigma_*$, depending on how it is distributed. We were unable to test this with the simulated observations described in this paper because the dust was not purely diffusely distributed.

When comparing spectral modelling methods, we found that using a Gauss-Hermite series to model the LOSVD yielded lower values of $f\sigma_*$, on average, compared with a pure Gaussian LOSVD when dust was included in the measurement. There was no systematic offset between the Gauss and Gauss-Hermite methods when dust was ignored.

Based on our findings, we can provide some general advice regarding comparisons between observationally-obtained measurements of σ_* (which is really $f\sigma_*$) and measurements of $m\sigma_*$ in the galaxy simulation literature:

- In systems that appear to be passively evolving (i.e., systems that clearly contain only one nucleus and little or no tidal debris), most measurements of $f\sigma_*$ are likely to fall within 20% of the $m\sigma_*$ measurement, regardless of the amount of dust present in the system. Even in extreme cases, $f\sigma_*$ should not differ from $m\sigma_*$ by more than 30% in such systems.
- Measurements of $f\sigma_*$ in passively evolving gas-rich systems with recent or ongoing star formation are most likely slightly positively offset with respect to $m\sigma_*$.
- If a large number of young stars (younger than ~ 200 Myr) in a passively evolving galaxy are not significantly obscured by dust, a measurement of $f\sigma_*$ is likely to be slightly negatively offset with respect to $m\sigma_*$.
- In systems that are actively merging (i.e., systems with two or more visible nuclei, more than one large disk structure, or an otherwise disturbed appearance), $f\sigma_*$ can differ from $m\sigma_*$ by more than 100% in extreme cases.

Finally, we note some important caveats to our findings:

- The numerical simulations that were used to create the galaxy snapshots were not perfect. Notably, the resolution limit was 25 pc, which means that the structure of the ISM was not resolved on scales of $\lesssim 25$ pc. This also means that the particles in the simulation and the subsequent Sunrise radiative transfer scheme represented entire stellar populations, with no sub-structure.
- We simulated a small region of the spectrum (5040 Å to 5430 Å). The results may differ somewhat in different regions of the spectrum. On the other hand, this region is commonly used in real observations of velocity dispersion, so it is a useful choice.
- Changing the dust grain model, dust-to-metal ratio, or total gas dust in the systems would have likely effected the details of the results. The magnitudes of the offsets would likely be different if any of these parameters were changed, however the direction of each offset is likely insensitive to to these variables.
- Our observational sample was quite small and it was not randomly chosen. It included some snapshots with highly unusual dynamics and specially-selected viewing directions. The extreme nature of the sample allowed us to identify exceptional cases which might not have been observed if the snapshots and viewing directions were chosen at random. However, one cannot compute robust statistics from this sample.

ACKNOWLEDGEMENTS

We thank Gillian Wilson and Desika Narayanan for their invaluable assistance. This work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number OCI-1053575. Financial support for this work was provided by NASA through a grant from the Space Telescope Science Institute (Program numbers AR-12626 and GO-11557), which is operated by the Association of Universities for Research in Astronomy, Incorporated, under NASA contract NAS5-26555.

REFERENCES

Baes M., Dejonghe H., 2000, MNRAS, 313, 153
Baes M., Dejonghe H., 2002, MNRAS, 335, 441
Bender R., Burstein D., Faber S. M., 1992, ApJ, 399, 462
Cappellari M., 2012, pPXF: Penalized Pixel-Fitting stellar kinematics extraction (ascl:1210.002)
Cappellari M., Emsellem E., 2004, PASP, 116, 138
Cox T. J., Dutta S. N., Di Matteo T., Hernquist L., Hopkins P. F., Robertson B., Springel V., 2006, ApJ, 650, 791
Davies R. L., Burstein D., Dressler A., Faber S. M., Lynden-Bell D., Terlevich R. J., Wegner G., 1987, ApJS, 64, 581
Dopita M. A., et al., 2005, ApJ, 619, 755
Dressler A., Lynden-Bell D., Burstein D., Davies R. L., Faber S. M., Terlevich R., Wegner G., 1987, ApJ, 313, 42
Ferrarese L., Merritt D., 2000, ApJ, 539, L9
Gebhardt K., et al., 2000, ApJ, 539, L13
Groves B., Dopita M. A., Sutherland R. S., Kewley L. J., Fischera J., Leitherer C., Brandl B., van Breugel W., 2008, ApJS, 176, 438
Johansson P. H., Burkert A., Naab T., 2009, ApJ, 707, L184
Jonsson P., 2006, MNRAS, 372, 2
Jonsson P., Primack J. R., 2010, Nature, 15, 509
Jonsson P., Groves B. A., Cox T. J., 2010, MNRAS, 403, 17
Leitherer C., et al., 1999, ApJS, 123, 3
Robertson B., Cox T. J., Hernquist L., Franx M., Hopkins P. F., Martini P., Springel V., 2006, ApJ, 641, 21
Rothberg B., Fischer J., 2010, ApJ, 712, 318
Rothberg B., Fischer J., Rodrigues M., Sanders D. B., 2013, ApJ, 767, 72
Springel V., 2005, MNRAS, 364, 1105
Stickley N. R., Canalizo G., 2012, ApJ, 747, 33
Stickley N. R., Canalizo G., 2014, ApJ, 786, 12
Tremaine S., et al., 2002, ApJ, 574, 740
Villumsen J. V., 1982, MNRAS, 199, 493
Weingartner J. C., Draine B. T., 2001, ApJ, 548, 296

APPENDIX A: APPENDIX: DESCRIPTIONS OF OBSERVATIONS

In this appendix, we describe each of the 21 individual observations (snapshot and camera position).

Snapshot 1

This snapshot was recorded at $t = 2.06$ Gyr—immediately before the climax of the second pass (the first pass occurred at 0.49 Gyr). We knew from the analysis performed in SC2 that the mass-weighted velocity dispersion reached its highest value during this stage of the merger. Consequently, this was a dynamically extreme system. The system contained two disks which were in the process of passing through one another. See Figure A1 for renderings of the system from the three camera positions and refer to the upper two panels of Figure 2 to see how this snapshot fits into the overall evolution of the merger.

Camera 1 was placed along a line of sight containing maximal dust attenuation. This line of sight fell approximately 27° away from the collision axis. Camera 2 was placed nearly perpendicular to the collision axis along a line of sight with minimal dust extinction. Camera 3 observed an intermediate configuration with a moderate amount of attenuation due to dust and a viewing angle 33° from the collision axis.

Snapshot 2

This snapshot was recorded at $t = 2.089$ Gyr. Dynamically, the system was still somewhat excited, having recently undergone a major collision. The system consisted of two clearly distinguishable, disturbed ellipsoids with extended disk debris.

Camera 1 was placed along a line of sight with a moderate degree of dust attenuation, 24° from the collision axis. Camera 2 was placed along a line of sight perpendicular to the collision axis with moderate attenuation. Camera 3 was placed approximately 5° from the collision axis. Attenuation was somewhat more significant in the Camera 3 observation than in the other two. A significant number of stars from both spheroids appeared in all slits of Cameras 1 and 3, whereas the slits used in the Camera 2 measurements primarily included stars belonging to only one of the ellipsoids. See Figure A2 for renderings.

Snapshot 3

This snapshot was recorded at $t = 2.181$ Gyr, which is approximately midway between the second and third passes of the merger process. The system consisted of two disturbed ellipsoidal galaxies. Both ellipsoids had nearly reached a dynamically stable state at this point. Enhanced star-formation, which was triggered during the second pass, was recently quenched by periods of quasar activity.

Cameras 1 and 2 were placed along lines of sight with moderate dust attenuation. Both sub-systems were clearly distinguishable from these camera positions. Camera 3 was placed along the line of sight connecting the two systems. Attenuation was more significant along this camera direction than along the first two camera directions. See Figure A3 for renderings.

Snapshot 4

This snapshot was recorded at $t = 2.299$ Gyr—the climax of the third pass. The two progenitor systems were fully superimposed; only one nucleus could be identified. Like Snapshot 1, this snapshot represented a highly unusual system, partially due to its elevated velocity dispersion, but also because the system hosted a weak nuclear starburst with a star formation rate of ~ 5 $M_\odot$ yr^{-1}.

Cameras 1, 2, and 3 were positioned 29°, 38°, and 40° from the collision axis. Dust attenuation was moderate in all three cases. See Figure A4 for renderings.

Figure A1. GSNAP-generated renderings of Snapshot 1 ($t = 2.060$ Gyr), viewed along the three camera directions that were examined. Each image represents a 100×100 kpc region. Green cross-hairs indicate the position of the centre of the slits. The size and orientation of the slits is shown in the upper-left corner, for reference. The first column shows the system with all stars and dust dust attenuation included. The second column shows the system without dust attenuation. The third column shows only the stars that formed during the simulation, attenuated by dust. The fourth column shows the image in the third column in the absence of dust attenuation. The brightness scaling in all images is identical.

Snapshot 5

This snapshot was recorded at $t = 2.401$ Gyr—the moment of nuclear coalescence. The majority of the merger's dynamical evolution was complete at this point. A nuclear starburst with a star formation rate of ~ 13 $M_\odot$ yr^{-1} was present. Major quasar activity began $\lesssim 5$ Myr after this snapshot was recorded, causing a sharp decrease in the star formation rate. Therefore, this system was likely similar to some quasar host galaxies.

Since the system was nearly isotropic dynamically as well as in terms of its dust distribution, the camera positions were chosen essentially at random. We used a large angular separation between viewing directions in order to prevent us from sampling the same regions of the simulation more than once. See Figure A5 for renderings of Snapshot 5.

Snapshot 6

This snapshot was recorded at $t = 2.690$ Gyr. The system exhibited shells and tidal debris in its outer regions. It also contained two non-coaxial, concentric nuclear disks with diameters of 0.3 kpc and 1.3 kpc. The inner disk was inclined 67° with respect to the outer disk.

Camera 1 was placed along a line of sight inclined $\sim 40°$ with respect to the direction defined by the intersection of the planes containing the two disks. This inclination angle was measured along a third plane that symmetrically bisected the system. In other words, the two disks were viewed from equal inclination angles. Camera 2 was placed such that both disks were viewed edge-on. Camera 3 was placed such that the larger disk was viewed edge-on. See Figure A6 for renderings.

Figure A2. The same as Figure A1, but for Snapshot 2 ($t = 2.089$ Gyr).

Snapshot 7

This snapshot was recorded at $t = 3.62$ Gyr—the end of the simulation. The system consisted of an elliptical galaxy containing three small nuclear disks with diameters of 9.0 kpc, 1.5 kpc, and 0.3 kpc. It may be more appropriate to refer to the largest disk, as a "ring," since it clearly exhibited outer and inner edges. See Figure A7 for renderings.

Camera 1 viewed the large ring face-on; the intermediate disk was inclined $\sim 20^\circ$ with respect to this line of sight. Camera 2 viewed the small inner disk edge-on. Camera 3 viewed the intermediate disk face-on. Cameras 2 and 3 respectively fell along lines of sight inclined $\sim 20^\circ$ and $\sim 23^\circ$ with respect to the edge of the large ring.

This paper has been typeset from a TEX/LATEX file prepared by the author.

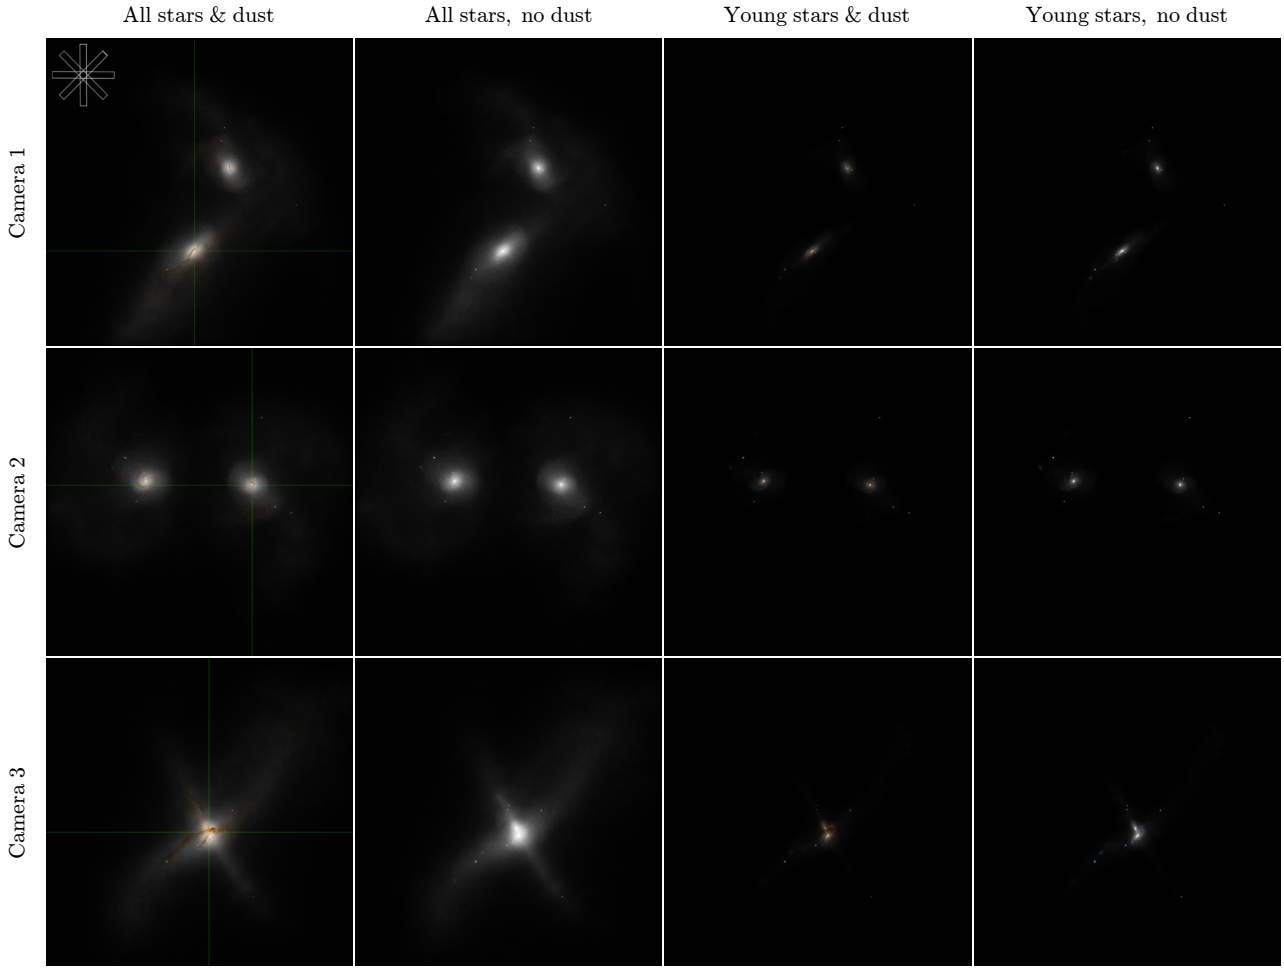

Figure A3. The same as Figure A1, but for Snapshot 3 ($t = 2.181$ Gyr).

Figure A4. The same as Figure A1, but for Snapshot 4 ($t = 2.299$ Gyr).

Figure A5. The same as Figure A1, but for Snapshot 5 ($t = 2.401$ Gyr).

Figure A6. The same as Figure A1, but for Snapshot 6 ($t = 2.690$ Gyr).

Figure A7. The same as Figure A1, but for Snapshot 7 ($t = 3.620$ Gyr).